

When Do AI Gains Become Broadly Shareable? A Policy Threshold for AI-Driven Automation

Aran Nayebi

Machine Learning Department and Neuroscience & Robotics Institutes,
School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA
anayebi@cs.cmu.edu

Abstract

AI-driven automation generates broad-based social benefit only if technical gains become visible, durable, and publicly claimable. We develop a policy-facing stress test by extending a standard task-automation growth model with an AI capability parameter that raises productivity on automatable tasks while holding the set of tasks fixed. The exercise is intentionally limited: it is not a forecast of AI timelines or a full welfare analysis, but a way to identify which institutions determine whether AI rents can support broad transfers. Calibrated to U.S. quantities, the model shows that capability alone is not decisive. Public capture, deployment costs, automation scope, and market structure jointly determine when AI gains become shareable. The main policy lesson is that moving from low to moderate public capture (33%) can substitute for substantial AI capability growth, while pushing capture further to full nationalization yields smaller gains, especially if deployment or safety costs are high. Competition policy also has distributional consequences: opening concentrated AI markets may improve fairness and resilience, but can reduce the rent pool unless alternative public-claim institutions are built. Cross-nationally, tax-heavy systems lower the needed AI capability threshold through stronger effective revenue collection, while Singaporean and Abu Dhabi-style public-asset models show that governments can also capture AI gains through ownership and investment returns rather than taxes alone. Our framework therefore identifies which levers governments can act on now to make future AI gains easier to measure, claim, and distribute broadly.

1 Introduction

AI governance debates often treat technical capability as the central variable: if AI becomes sufficiently powerful, then society can decide how to distribute the gains (Dafoe 2018; OpenAI 2018, 2023; O’Keefe et al. 2020; Anderljung et al. 2023). This paper starts from the *opposite* premise. Capability growth matters, but it does not by itself determine whether AI-driven productivity gains become broad public benefit, concentrated private rent, lower prices, worker augmentation, or offshored profit. That outcome depends on institutions: taxation, public ownership, anti-avoidance capacity, market concentration, procurement rules, and the design of AI deployment. Foundation models already sit

inside cloud contracts, chip supply chains, enterprise software, labor processes, and public-facing services; the resulting rents are shaped by technical performance, but also by market concentration, ownership structure, and tax enforcement across the AI value chain (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b,a).

We develop a policy benchmark for this institutional problem. Rather than asking only how “transformative” AI systems may become, our benchmark asks: **Under what institutional conditions do AI-generated rents become sufficiently observable, capturable, and durable to support broad-based social benefit?** Throughout, “AI rents” denotes AI-attributable capital income available for public capture; Proposition 2 separately adds pure economic profit created by market power. Our benchmark links AI capability to the public capture of automation rents, but it is not tied to any single redistributive instrument. A broad transfer can be implemented as an AI dividend, a sovereign wealth dividend, a refundable tax credit, a negative income tax, or a universal basic income. Therefore, what we consider here is the *size* of the broad-based transfer benchmark, and leave the precise administrative form used to distribute it to be flexibly chosen.

Specifically, our contribution is to extend traditional task-automation growth models with an AI capability parameter and use the resulting threshold condition to connect technical progress to institutional choices about public capture, market structure, and broad-based benefit. Task automation and capital deepening have long been central in macroeconomics (Solow 1956, 1957; Baumol 1967; Zeira 1998; Acemoglu and Restrepo 2018). Recent AI-growth models and reviews ask how automation, new tasks, or explosive growth may change aggregate output (Aghion, Jones, and Jones 2017; Erdil and Besiroglu 2023). Our question is complementary: given a stream of AI-generated rents, we ask when existing institutions can convert those rents into broad social benefit. To make the benchmark demanding and interpretable, we calibrate it to the standardized 11%-of-GDP UBI (Universal Basic Income)-sized transfer (UBI Center 2019; U.S. Bureau of Economic Analysis 2024c), while emphasizing that UBI is a yardstick rather than the paper’s normative endpoint. To isolate the institutional problem, we hold the task frontier fixed and do not assume that AI cre-

ates new jobs, new tasks, or compensating labor demand. This makes the model a conservative policy *stress test* rather than a speculative forecast.

Within that environment, we derive a *closed-form* capability threshold linking AI productivity on automatable tasks to effective public capture, operating and compliance costs, automation intensity, and market structure. This reframing matters because these quantities correspond to *observable policy levers*. Public capture depends on tax enforcement, withholding rules, public equity, licensing, and anti-avoidance capacity (U.S. Government Accountability Office 2022; Organisation for Economic Co-operation and Development 2025c,d, 2026). Operating costs depend on compute, inference, safety, energy, and compliance burdens. Market structure affects whether AI rents appear as competitive surplus, concentrated profit, or bargaining power within the AI value chain (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b,a). The benchmark therefore does not ask only whether AI becomes “transformative”; rather, it asks whether governments can observe, capture, and govern the rents generated by AI deployment.

Using current U.S. macroeconomic quantities, our central calibration implies that supporting an 11%-of-GDP broad transfer would require AI to reach roughly 5-7 times today’s automation productivity on automatable tasks. Raising effective public capture from roughly 15% to 33% most significantly lowers the threshold to about 3.2 $\times$, with diminishing returns when AI companies are fully nationalized. Market concentration lowers the funding threshold by expanding rents, but this is a fiscal result rather than a welfare endorsement. The same concentration that expands the capturable rent pool can also create exclusion, lock-in, dependence, and bottlenecks in AI infrastructure and downstream markets (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b,a). The policy implication is that capability policy, competition policy, and distribution policy must be designed *together*.

We then use our benchmark comparatively, interpreting it through U.S., Nordic, Japanese, Singaporean, and Abu Dhabi-style institutional baselines. Recent comparative governance work emphasizes that generative-AI governance principles vary across regions, that compute-based governance capacity is unevenly distributed, and that systematic monitoring infrastructure remains underdeveloped (Luna et al. 2024; Lehdonvirta, Wu, and Hawkins 2024; Arnold et al. 2024). Our benchmark complements that literature by translating institutional differences into capability thresholds for broad-based benefit. Countries with different levels of effective public capture, anti-avoidance capacity, public-asset ownership, and value-chain position do not face the same boundary between privately captured AI rents and publicly shareable AI gains.

The result is *not* a prediction of when AI will fund a transfer program, nor a decision rule that mechanically determines policy. It is a diagnostic tool for asking which institutional constraints make broad-based AI benefit more or less fiscally plausible. Current empirical evidence sug-

Lever	Meaning in our model	Observable proxy
γ_t	AI productivity on automated tasks	Task-duration success rates, deployment benchmarks, and economically meaningful evaluations of autonomous task completion (Model Evaluation and Threat Research 2025).
$\bar{\alpha}$	Share of tasks that can be automated	Task and occupation exposure measures from workforce surveys, firm adoption data, and task-level automation studies (World Economic Forum 2023).
Θ	Effective public capture of AI rents	Effective corporate tax collection, withholding regimes, public equity claims, and anti-avoidance enforcement outcomes (U.S. Government Accountability Office 2022; Organisation for Economic Co-operation and Development 2025c,d).
c	Operating and compliance cost share	Compute, inference, safety, energy, and regulatory expenditures as a share of gross AI revenue or rent.
ψ_t	Augmentation before full automation	Worker-level productivity gains, time savings, and changes in task mix under AI assistance (Brynjolfsson, Li, and Raymond 2023; Dillon et al. 2025; Anthropic 2026).

Table 1: Governance levers can be observed rather than treated as abstractions. Each model parameter can be tied to a concrete measurement or institutional proxy.

gests that many AI deployments still augment workers rather than replace them outright (Brynjolfsson, Li, and Raymond 2023; Dillon et al. 2025; Anthropic 2026). Recent work from METR also suggests that governance-relevant capability measurement should be task-based and deployment-relevant, rather than limited to static model-evaluation scores (Model Evaluation and Threat Research 2025). We therefore present our framework as a policy test for AI-driven automation, connecting fiscal feasibility to capability monitoring, anti-avoidance enforcement, public-asset strategies, procurement, competition policy, and labor augmentation.

2 Observable Governance Levers

Before turning to the formal model, Table 1 maps our key parameters to observables and institutional settings. We place this map first because the model is meant to organize and *prioritize* real policy levers, not reduce governance to a single capability score.

Table 1 also clarifies why we interpret AI as a value-chain and governance problem rather than a pure production-function problem. Capability gains do not automatically become public benefit. They may be dissipated by compute, safety, or compliance costs (c); retained privately when public capture (Θ) is weak; or shifted across jurisdictions if tax

and reporting rules are weak. Conversely, AI systems do not need to have automated the whole economy for governance questions to become pressing. Even partial automation ($\bar{\alpha}$) of high-value tasks can generate meaningful rents, while ψ_t separately tracks gains from complementing workers before full automation (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b).

For monitoring, γ_t should be estimated sector by sector as quality-adjusted productivity on a fixed basket of automatable tasks relative to that sector's pre-AI automation baseline. An economy-wide index should use fixed value-added weights and report cross-sector dispersion; task-duration benchmarks are leading indicators rather than direct estimates of economic productivity. Effective public capture should be estimated as net AI-attributable taxes, withholding, royalties, licensing revenue, and public-asset returns, less AI-specific subsidies, divided by the corresponding AI-attributable capital-income base. The statutory rates used below are therefore scenario anchors, not estimates of national Θ .

Seen this way, Table 1 is also a map of institutional responsibility. Model developers and cloud providers influence c through compute efficiency, safety practices, and compliance architecture. Downstream deployers influence $\bar{\alpha}$ and ψ_t through workflow redesign, staffing choices, and adoption sequencing. Tax and finance authorities influence Θ through reporting, withholding, and anti-avoidance enforcement. The same deployment decision can therefore affect both the distribution of benefits and the amount of rent available to share in the first place (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b,a).

Finally, beginning with this table helps us keep two transition channels conceptually separate. The first is *augmentation*, which we represent with ψ_t and which is already visible in current deployments. The second is *automation rent extraction*, which depends on both $\bar{\alpha}$, the share of tasks exposed to automation, and γ_t , the productivity of AI on those automated tasks. The same solvency benchmark applies across both these channels, with augmentation shaping near-term output and automation rents determining the longer-run funding base.

2.1 A Large UBI Benchmark as a Useful Common Yardstick

A calibration needs a common fiscal scale. We use an 11%-of-GDP UBI-sized transfer because it is demanding enough to be policy-relevant, legible across disciplines, and administratively neutral in the model. It is a yardstick rather than a recommendation: the same aggregate funding condition can support AI dividends, refundable credits, wage supplements, public services, or other broad-benefit instruments.

This separation also clarifies our normative scope. The threshold is a positive fiscal-feasibility result conditional on the premise that some AI gains should be shared broadly; it does not identify the welfare-optimal transfer, delivery mechanism, or recipients. Readers therefore need not endorse UBI to use the benchmark.

Equation (8) also makes smaller targets transparent. Holding other parameters fixed, a transfer share $b = B/Y_t$ rescales as

$$\gamma_t^*(b) = \gamma_t^*(0.11) \left(\frac{b}{0.11} \right)^\sigma. \quad (1)$$

At the central $\sigma = 0.66$, a 5%-of-GDP target lowers the current-capture benchmark from $5.6\times$ to about $3.3\times$, and the one-third-capture benchmark from $3.2\times$ to about $1.9\times$. Thus targeted instruments are not merely qualitatively easier to fund; their thresholds can be read directly from the same model.

3 Defining Our Conservative Automation Benchmark: Solow-Zeira Growth Model with an AI Capability Parameter

3.1 Holding the Task Frontier Fixed Gives a Policy Stress Test

We model GDP Y_t with a constant-elasticity-of-substitution (CES) production function, which aggregates task-level goods X_{it} into total output. In policy terms, this lets us represent the economy as a collection of tasks that differ in whether output is produced by labor or by automated capital. Specifically, at date t a unit-measure continuum of tasks $i \in [0, 1]$ is aggregated through

$$Y_t = A_t \left(\int_0^1 X_{it}^\rho di \right)^{1/\rho}, \quad \rho = \frac{\sigma - 1}{\sigma} < 0, \quad (2)$$

where $A_t = A_0 e^{gt}$ is Hicks-neutral progress with growth rate $g \in \mathbb{R}$ captures generic efficiency improvements (e.g. Moore's law, better chips, cheaper energy, learning curves) that raise output even when no new jobs are created. $\sigma \in (0, 1)$ is the elasticity of substitution across tasks, where a lower σ means that tasks are harder to substitute for one another, so bottlenecks in non-automated work continue to matter even when automated tasks become much more productive. As in Zeira (1998), each task is produced by either automated capital or labor:

$$X_{it} = \begin{cases} K_{it} & \text{if task } i \text{ is automated at time } t, \\ L_{it} & \text{otherwise.} \end{cases} \quad (3)$$

Here K_{it} denotes capital services used for an automated task, including machines, software, compute, and AI-enabled systems, while L_{it} denotes labor on a non-automated task. The production block is technology-agnostic: software-based generative AI, embodied robotics, and hybrid systems can all enter through deployment-specific estimates of γ_t and $\bar{\alpha}$.

Specifically, we fix a fraction $\bar{\alpha} \in (0, 1]$ of tasks as automatable. As a result, we do not need to assume full automation; rather, $\bar{\alpha}$ is itself a governance-relevant lever: it captures how much of the economy can, in principle, be shifted into the automated block. Holding it fixed gives a conservative benchmark, which allows us to ask whether broad-based transfers remain feasible even if AI automates only the currently automatable share and no new tasks emerge.

With uniform allocation across the automated and non-automated blocks, i.e. $K_{it} \equiv K_t/\bar{\alpha}$ for automated tasks and $L_{it} \equiv L$ for non-automated tasks, output simplifies to

$$Y_t = A_t (\bar{\alpha}^{1-\rho} K_t^\rho + (1 - \bar{\alpha})^{1-\rho} L^\rho)^{1/\rho}. \quad (4)$$

Here K_t denotes aggregate capital services allocated to the automated block, while L denotes labor allocated to the non-automated block. The two terms inside the production function therefore correspond to the automated and non-automated parts of the economy. This simplified expression is useful because it makes the central governance tradeoff explicit: broad-based benefit depends not only on how capable AI becomes, but also on how much of the economy is exposed to automation and how much of the resulting capital income can be publicly captured.

3.2 Capability Should Scale Automated-Task Productivity, Not Raw Capital

We model AI capability γ_t as *raising the productivity of automated tasks*. The parameter γ_t is measured relative to pre-AI automation, so $\gamma_t = 1$ is the baseline and $\gamma_t > 1$ means that AI makes already-automated tasks more productive. Rather than treating AI as simply adding more raw capital, we let capability scale the effective contribution of the automatable block:

$$\bar{\alpha}^{1-\rho} K_t^\rho \mapsto (\gamma_t \bar{\alpha})^{1-\rho} K_t^\rho = \gamma_t^{1-\rho} \bar{\alpha}^{1-\rho} K_t^\rho.$$

The production function therefore becomes

$$Y_t = A_t \left(\gamma_t^{1-\rho} \bar{\alpha}^{1-\rho} K_t^\rho + (1 - \bar{\alpha})^{1-\rho} L^\rho \right)^{1/\rho}, \quad \gamma_t \geq 1. \quad (5)$$

This specification has a direct policy interpretation: AI capability does not expand the literal set of automatable tasks; that role is played by $\bar{\alpha}$; instead, γ_t measures how much more economically important the automated block becomes once AI improves the productivity of tasks already exposed to automation. This distinction is useful because the governance question is not only how many tasks can be automated, but also how much rent those automated tasks can generate once AI systems become more capable.

The choice to scale the automated block rather than raw capital is also economically motivated. If we multiplied capital directly by γ_t , the capability boost would enter as γ_t^ρ and would be mechanically attenuated when $\rho < 0$. That would make higher AI capability look less important precisely in the case where tasks are hard to substitute for one another. For our purposes, that is the wrong object: we want γ_t to capture the productivity of AI on automatable tasks, not simply the quantity of capital used in them.

The condition $\rho < 0$, equivalently $\sigma < 1$, means that tasks are gross complements. In policy terms, this captures the Baumol “cost-disease” effect (Baumol 1967): even if AI makes automated tasks much cheaper or more productive, non-automated tasks such as care, oversight, judgment, service delivery, or other labor-intensive activities can remain economically important. As automated tasks become more productive, the economy does not necessarily become all automation all the way down. Instead, relative scarcity shifts

toward the slower-moving non-automated block. This keeps the model from assuming that AI capability alone makes labor irrelevant, while still allowing higher γ_t to expand the rent pool generated by automation.

Finally, we note that restricting attention to $\gamma_t \geq 1$ is without loss of generality. Values below one would mean that AI makes the automated block less productive than pre-AI automation. In the model, such a case can be absorbed by redefining the effective automatable share as $\bar{\alpha}' = \gamma_t \bar{\alpha}$, so it adds no distinct economics. We therefore treat $\gamma_t = 1$ as the pre-AI automation baseline and $\gamma_t > 1$ as any additional productivity boost delivered by AI, if any.

3.3 Effective Public Capture is a Key Fiscal Lever

To determine how much rent can be shared, we also need a simple way to summarize the economy’s capital base. We assume that a constant share s of output is saved and invested, while existing capital depreciates at rate δ via the standard Solow capital-accumulation difference equation (Solow 1956, 1957):

$$K_{t+1} = sY_t + (1 - \delta)K_t. \quad (6)$$

This standard accumulation equation says that next period’s capital equals new investment plus the undepreciated part of the current capital stock.

Once AI capability γ_t stabilizes, or grows at any constant factor $G_\gamma := \gamma_{t+1}/\gamma_t > 0$ while the capital stock adjusts at $G_K := K_{t+1}/K_t = G_\gamma^{1/(1-\sigma)}$, so that $Y_{t+1} = e^g Y_t$ in either case, this linear difference equation forces the capital-output ratio to a Solow limit of

$$\frac{K_t}{Y_t} \xrightarrow{t \rightarrow \infty} \bar{\kappa} = \frac{s}{e^g - 1 + \delta} \stackrel{|g| \ll 1}{\approx} \frac{s}{g + \delta}. \quad (7)$$

We use this ratio to express the rent-generating automated capital base relative to the size of the economy. This keeps the threshold in terms of observable macro quantities rather than requiring a separate forecast of the total capital stock.

We model the public sector as capturing a share $\Theta \in (0, 1]$ of AI rents, net of an operating and compliance cost share $c \in [0, 1)$. We interpret Θ broadly as *effective* public capture through taxation, withholding, anti-avoidance enforcement, public ownership, or equivalent claims on AI-generated rents. This broader interpretation is important: in the comparative section below, we use real-world tax regimes as stylized baselines for Θ , but we also distinguish tax-heavy systems from public-asset models whose effective public claim can exceed what headline tax rates alone suggest.

4 Identifying Tractable Governance Thresholds

Our model, defined in §3, yields a transparent benchmark for the capability required to finance a transfer benchmark B :

Proposition 1 (Capability threshold). *In the economy described by (5)–(7), a transfer benchmark B is fiscally sus-*

tainable in period t if and only if

$$\gamma_t \geq \gamma_t^* = \left(\frac{B/Y_t}{\Theta(1-c)\bar{\alpha}^{1-\rho}A_t^\rho\bar{\kappa}^\rho} \right)^\sigma, \quad \sigma = \frac{1}{1-\rho} > 0. \quad (8)$$

Proofs of all propositions and corollaries are available in the full arXiv version: <https://arxiv.org/abs/2505.18687>.

Proposition 1 is the benchmark around which the rest of our analysis turns. It says that the AI capability threshold γ_t^* is *multiplicatively* shaped by public capture (Θ), AI training and safety costs (c), automation intensity ($\bar{\alpha}$), and the capital-output ratio ($\bar{\kappa}$). In particular, AI capability and public capture are not substitutes in any simple political sense. A society can lower the required capability by raising capture, or lower the required capture by waiting for more capable AI, but the trade-off is dictated by (8).

Corollary 1 (Comparative statics). *For the threshold*

$$\gamma_t^* = \left(\frac{B/Y_t}{\Theta(1-c)\bar{\alpha}^{1-\rho}A_t^\rho\bar{\kappa}^\rho} \right)^\sigma, \quad \bar{\kappa} = \frac{s}{e^g - 1 + \delta},$$

the first-order effects are

$$\begin{aligned} \frac{\partial \gamma_t^*}{\partial \Theta} &= -\frac{\sigma \gamma_t^*}{\Theta} < 0, & \frac{\partial \gamma_t^*}{\partial c} &= \frac{\sigma \gamma_t^*}{1-c} > 0, \\ \frac{\partial \gamma_t^*}{\partial \bar{\alpha}} &= -\frac{\gamma_t^*}{\bar{\alpha}} < 0, & \frac{\partial \gamma_t^*}{\partial \bar{\kappa}} &= -\frac{\sigma \rho \gamma_t^*}{\bar{\kappa}} > 0, \\ \frac{\partial \gamma_t^*}{\partial s} &= -\frac{\sigma \rho \gamma_t^*}{s} > 0, & \frac{\partial \gamma_t^*}{\partial \delta} &= \frac{(\sigma-1)\gamma_t^*}{e^g - 1 + \delta} < 0. \end{aligned} \quad (9)$$

The comparative statics give insight into governance: More public capture (Θ) lowers the required AI capability for broad-based benefits. Higher operating or safety costs (c) raise it. Broader automation $\bar{\alpha}$ lowers it. A higher capital-output ratio $\bar{\kappa}$ raises the threshold because the automated-capital rent share enters as $\bar{\kappa}^\rho$ with $\rho < 0$: at a fixed capability level, additional capital reduces the marginal rent share of the automated block. The saving and depreciation effects follow through this same capital-output channel: higher saving s raises $\bar{\kappa}$ and therefore raises the required capability, while higher depreciation δ lowers $\bar{\kappa}$ and therefore lowers it.

Proposition 2 (Funding threshold under oligopoly). *Suppose the AI-capital sector is supplied by Cournot firms with conduct parameter $\theta = \sum_i s_i^2 \in [0, 1]$ and demand elasticity $\varepsilon > 0$, so pure profit accounts for the share θ/ε of output. If the public sector captures that pure-profit component in addition to the rent share in Proposition 1, then the transfer benchmark is sustainable if and only if*

$$\gamma_{\text{oligo},t}^* = \left(\frac{B/Y_t}{\Theta(1-c)} \frac{1}{\bar{\alpha}^{1-\rho}A_t^\rho\bar{\kappa}^\rho} - \frac{\theta/\varepsilon}{\bar{\alpha}^{1-\rho}A_t^\rho\bar{\kappa}^\rho} \right)^\sigma. \quad (10)$$

Whenever $\theta > 0$, $\gamma_{\text{oligo},t}^ < \gamma_t^*$.*

Equation (10) should *not* be read as saying that concentrated AI markets are good. It says only that, if a small number of firms earn large AI rents, there is a larger pool of money that governments could in principle capture and redistribute, thereby making a broad transfer easier to fund. But the same concentration can also make AI markets less fair and less resilient by increasing lock-in, limiting entry, raising switching costs, and making governments and users dependent on a few suppliers.

Corollary 2 (Higher public capture lowers the necessary AI capability threshold). *If $\Theta_2 > \Theta_1$, then*

$$\gamma_{2,t}^* < \gamma_{1,t}^* \quad \text{and} \quad \gamma_{1,t}^* - \gamma_{2,t}^* = C_t (\Theta_1^{-\sigma} - \Theta_2^{-\sigma}) > 0, \quad (11)$$

where

$$C_t = \left(\frac{B/Y_t}{(1-c)\bar{\alpha}^{1-\rho}A_t^\rho\bar{\kappa}^\rho} \right)^\sigma. \quad (12)$$

Corollary 2 is especially useful for comparative analysis. It tells us that cross-national differences in effective public capture map directly into different capability thresholds. That is the compact formal statement behind the comparative discussion below: countries do not need the same level of AI capability to support broad-based benefit if their institutions turn rents into public revenue at different rates, as we now analyze below.

5 Calibration and Sensitivity Analyses

5.1 Moderate Public Capture Substantially Lowers the Needed Capability

We start by calibrating the benchmark using current U.S. quantities where possible. The transfer benchmark is $B/Y_t \approx 0.11$, corresponding to a \$12,000 annual transfer for adults relative to 2024 U.S. GDP (UBI Center 2019; U.S. Bureau of Economic Analysis 2024c). We set public capture $\Theta = 0.145$ using GAO evidence on effective federal corporate income taxation (U.S. Government Accountability Office 2022). We use $\bar{\alpha} = 0.42$ based on automation-related task shares reported by the World Economic Forum (World Economic Forum 2023). For substitution, we use the U.S. capital-labor elasticity range summarized by Knoblach, Roessler, and Zwerschke (2020), centered at $\sigma = 0.66$. We set $s = 0.22$ from U.S. gross capital formation (World Bank 2025), $g = 0.011$ from CBO productivity projections (Congressional Budget Office 2024), and $\delta = 0.056$ from BEA fixed-asset data (U.S. Bureau of Economic Analysis 2024a,b). Because frontier AI operating and compliance costs remain highly uncertain, we treat c as a policy-sensitive range and use $c = 0.60$ as our central benchmark, conservatively informed by third-party estimates of OpenAI's gross margins and inference costs (Sacra 2025).

That last choice is worth underscoring. We do not want the benchmark to smuggle in a false precision about operating costs. Compute, inference, safety, energy, evaluation, and compliance expenditures may all change quickly as AI diffuses into more sectors. In other words, c is not just a nuisance parameter; it is one of the ways in which the institutional and technical shape of deployment enters the fiscal question. For that reason, we plot sensitivity bands in Figure 1 rather than a single point estimate.

Figure 1 summarizes the first-order benchmark. Under current public capture and benchmark costs, the threshold is about $5.6\times$ today's automation productivity on automatable tasks. Moving across plausible operating-cost values produces a range of roughly $4.8\times$ to $7.6\times$. Varying σ across the Knoblach, Roessler, and Zwerschke (2020) range yields a threshold between about $4.9\times$ and $6.4\times$. Raising public capture from about 15% to one-third lowers the threshold to




Figure 1: **Moderate public capture can substitute for substantial capability growth.** Thick segments show the threshold range induced by operating-cost uncertainty, thin inner segments show the elasticity range, and dots show the benchmark point estimate. Under current U.S. capture, the required AI capability threshold is roughly $5.6\times$ today’s automation productivity; under one-third public capture it falls to roughly $3.2\times$.

roughly $3.2\times$, with a broader sensitivity range of roughly $2.4\times$ to $4.4\times$ once elasticity and cost uncertainty are included. Institutional reform therefore changes the calibrated boundary by several multiples of today’s automation productivity.

The policy significance of Figure 1 is not that it predicts a date on which AI will inevitably cross the line. We intentionally avoid that framing here. The point is instead that the relevant threshold is not obviously science-fictional in scale. In our benchmark, institutions move the threshold by roughly the same order of magnitude as many speculative arguments about capability growth. That is precisely why governance choices matter *earlier* rather than later.

5.2 Market Concentration Can Lower the Funding Threshold While Worsening Welfare

Figure 2 illustrates Proposition 2. Imposing the symmetric-Cournot benchmark $\theta = 1/m$ makes the policy intuition transparent: most of the change happens when markets move from monopoly to a small oligopoly, not when they move from ten firms to twenty-five. The current AI sector sits much closer to concentrated oligopoly than to perfect competition in several layers of the stack, from specialized chips and cloud infrastructure to access pathways for frontier models (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b). In our benchmark, that concentration lowers the *funding* threshold because it creates larger profits. But it simultaneously raises governance concerns about dependency, foreclosure, and bargaining power within the AI value chain. Proposition 2 establishes no welfare ranking: consumer surplus, innovation, worker power, resilience, and democratic accountability remain outside its fiscal comparison.

This is a genuinely double-edged result for AI governance. If regulators successfully open markets and reduce rents, then broad-based benefit does not become less desirable; it simply has to be financed through more robust

public-capture instruments rather than by relying on concentrated markups. The right lesson is therefore not “monopoly is good,” but that competition governance and distributional governance cannot be analyzed in isolation.

5.3 Moving from Low to Moderate Capture Matters More Than Chasing Full Nationalization

Figure 3 shows that public capture is one of the strongest institutional levers in the model, but also that its gains are *not* linear. Starting from the U.S. benchmark, raising effective capture lowers the threshold sharply at first and then more slowly. The one-third benchmark is therefore attractive not because it is magically optimal, but because it captures the steep part of the curve. Past that point, additional capture still helps, but the marginal reduction in γ^* is noticeably smaller, especially when c is high.

That curvature is not incidental. Because $\gamma^* = C_t \Theta^{-\sigma}$ and $\sigma \in (0, 1)$, the largest threshold reductions occur when a country moves from very low capture toward moderate capture; later increments still help, but less dramatically. This is the formal reason the one-third benchmark matters so much in our calibration and why moving the threshold does not require full nationalization or anything close to it. Once operating and compliance costs are high, the practical challenge is not simply to raise headline rates, but to raise *effective* capture while keeping administrative drag and avoidance manageable.

This is where Corollary 2 becomes useful for comparative policy analysis. Cross-national tax and ownership regimes provide a rough sense of how different states sit at different points on the curve, even though statutory tax rates do *not* map one-to-one into Θ , as it can depend on tax bases, deductions, enforcement, profit shifting, withholding rules, public ownership, and other claims on AI-generated value. Sweden’s official corporation tax rate is 20.6% (Swedish Agency for Economic and Regional Growth 2025), while Norway’s




Figure 2: **Competition policy and distribution policy have to be designed together.** The solid curve shows the symmetric-Cournot threshold when conduct scales as $\theta = 1/m$ and the absolute demand elasticity is fixed at the UK Department for Science, Innovation and Technology (2023) midpoint of $|\varepsilon| = 1$; the dashed line marks the competitive benchmark. Moving from monopoly to a small oligopoly raises the required AI capability threshold sharply, while additional entry beyond that has smaller effects. The shaded region marks where competition reduces AI firms’ excess profits most quickly.

business general-income rate is 22% (The Norwegian Tax Administration 2025). Singapore’s flat corporate income tax rate is 17% (Inland Revenue Authority of Singapore 2026), and the UAE’s federal corporate tax rate is 9% above the relevant threshold (Ministry of Finance - United Arab Emirates 2026). For Japan, JETRO reports a 29.74% effective corporate tax rate for larger enterprises on a standard-rate basis (Japan External Trade Organization 2025). Using those values only as stylized baselines for effective capture, the implied benchmark threshold is roughly $5.6\times$ under the current U.S. baseline, $5.0\times$ under Singapore’s headline tax rate, $4.4\times$ under a Swedish-style rate, $4.2\times$ under a Norwegian-style rate, $3.5\times$ under the Japanese large-firm rate, $3.2\times$ at the one-third benchmark, and $7.7\times$ under the UAE’s tax-only federal baseline.

Table 2 should therefore be read as a comparison of jurisdictional fiscal capacity: countries with stronger institutions for turning concentrated profits into public revenue need less AI capability to support the same transfer benchmark. It is not a claim that the state hosting or taxing an AI firm has the sole moral entitlement to globally produced value.

Singapore and Abu Dhabi sharpen that comparative reading because both illustrate public-asset pathways that a tax-only proxy understates. Singapore’s Net Investment Returns Contribution allows the government to spend up to 50% of expected long-run returns on reserves invested by GIC, MAS, and Temasek, which means public capture can occur through public assets as well as through ordinary corporate taxation (Ministry of Finance Singapore 2025; GIC Private Limited 2026). Abu Dhabi combines the UAE’s low federal corporate tax with large sovereign investors such as ADIA, Mubadala, and ADQ, which are explicitly tasked with sus-

Illustrative baseline	Θ used	γ^* at $c = 0.60$
Current U.S. effective federal corporate tax	0.145	5.6
Singapore statutory corporate tax	0.170	5.0
Sweden statutory corporate tax	0.206	4.4
Norway business/general-income rate	0.220	4.2
Japan effective corporate tax rate (larger enterprises)	0.297	3.5
One-third public-capture benchmark ($\Theta = 0.33$)	0.330	3.2
UAE federal corporate tax (tax-only baseline)	0.090	7.7

Table 2: **Tax-heavy and public-asset states offer different routes to effective public capture.** These values use published tax rates as stylized reference points rather than direct estimates of effective AI-rent capture.

taining long-run prosperity, generating returns for the government, and diversifying the economy (Abu Dhabi Investment Authority 2026; Mubadala Investment Company 2026; ADQ 2026). The point is not that these institutions can be copied wholesale or that they already solve AI distribution. It is that low headline taxation and low effective public claim are not the same thing.

That distinction matters for current industrial-policy debates. Many proposals for AI dividends, public co-investment, sovereign or social wealth funds, data-center community-benefit agreements, and public-interest licensing are trying to move only three objects in our model: raise




Figure 3: **The biggest fiscal gains come from moving from low to moderate public capture.** The curves plot the required AI capability threshold as effective public capture rises under three operating-cost assumptions ($c = 0.50, 0.60, 0.75$). Axis markers indicate stylized comparison points for the UAE, United States, Singapore, a Nordic tax band, Japan, and the one-third benchmark. Because Θ is broader than statutory taxation alone, Singaporean and Abu Dhabi-style public-asset institutions should be read as tax-only lower bounds rather than complete estimates of public claim on AI rents.

Θ , keep c from being shifted onto the public, or increase visible worker benefit during the augmentation phase. Our framework validates that instinct, but it also clarifies the constraint: these ideas matter only if they produce durable public claims on rent without adding so much leakage or administrative drag that the gains are dissipated.

The comparative lesson is also about sequencing. Jurisdictions starting near the U.S. baseline cannot count on capability growth alone if they want broad-based benefit in a politically relevant time frame; they need to move on capture and enforcement earlier. Jurisdictions already closer to the one-third benchmark ($\Theta = 0.33$) may get more from improving measurement, competition governance, and transition support, because additional gains from higher capture arrive more slowly once they are already on the flatter part of Figure 3. The same global technology trend therefore implies different immediate priorities across states.

That observation matters especially because AI rents are likely to be internationally mobile. Cross-border payments, licensing arrangements, cloud contracts, and intellectual-property location choices all affect where rents appear to arise and where they are taxed. OECD evidence on withholding taxes, treaty networks, and minimum-tax coordination exists precisely because firms can shift profits when public capture is fragmented or poorly coordinated (Organisation for Economic Co-operation and Development 2025c,d, 2026). In an AI context, this means that public-capture policy is partly a domestic question and partly a co-ordination question. If countries compete to attract compute clusters or model firms by offering weak capture, the effective Θ relevant for broad-based benefit can stay low even as

aggregate rents rise.

Comparative perspective also means recognizing that jurisdictions do not occupy the same place in the AI value chain. A country that hosts model labs, cloud regions, or valuable downstream platforms can often capture profits more easily than a country that mainly imports AI services. For the latter, the central issue is often less which headline rate to choose than whether source-based withholding, reporting obligations, and co-ordinated audit rules prevent value generated in domestic deployment from being booked elsewhere. Cross-national inequality therefore enters our benchmark through institutions as much as through capability: the same γ_t can produce very different social outcomes depending on whether a state has the capacity to locate, measure, and claim part of the profit stream (Organisation for Economic Co-operation and Development 2025d, 2026).

Recent comparative governance scholarship reaches a similar conclusion. Luna et al. show that governance processes and principles for generative AI already differ meaningfully across regions, while Lehdonvirta, Wu, and Hawkins show that the geography of AI compute creates unequal opportunities for states to govern AI systems through territorial control of infrastructure (Luna et al. 2024; Lehdonvirta, Wu, and Hawkins 2024). Our comparative thresholds translate those institutional asymmetries into a more concrete question: how much capability must a country wait for if it cannot already locate, tax, or own a meaningful share of AI rents?

Real-world public-rent institutions also illustrate why the model's collective-benefit channel is not merely hypotheti-

cal. Norway's Government Pension Fund Global is financed by state petroleum revenues and invests those proceeds for long-run public benefit (Norwegian Ministry of Finance 2026). Petroleum rents are obviously not AI rents, and we do not mean to collapse the analogy. But the institutional lesson is important: concentrated economic rents can be converted into broad public assets through durable state capacity rather than through ad hoc redistribution alone. When public-asset pathways are used, however, governance quality matters too. The Santiago Principles exist precisely because sovereign wealth funds need transparent mandates, accountability, and risk-management discipline if they are to convert concentrated rents into legitimate public benefit (International Forum of Sovereign Wealth Funds 2026). That helps explain why our model treats UBI as one benchmark inside a broader family of dispersion mechanisms rather than as the only available endpoint.

The practical implication is that distribution mechanisms should be matched to institutional maturity. Some countries may be better placed to rely on general taxation and existing transfer systems; others may find royalty-style mechanisms, sectoral levies, or publicly governed funds easier to administer. Our benchmark does not adjudicate among those designs. It clarifies that they are all attempts to move the same underlying lever: effective public capture of AI rents.

6 Practical Policy Takeaways

Our benchmark is most useful as a policy checklist. Broad-based benefit depends not only on how capable AI systems become, but also on whether governments can capture part of the resulting rent, keep deployment costs visible, preserve contestability where it matters, and make near-term augmentation beneficial enough to sustain legitimacy. These are not separate debates: they map onto public capture (Θ), AI safety and operating costs (c), market concentration (θ), demand elasticity (ε), the share of tasks automated ($\bar{\alpha}$), and labor augmentation (ψ_t), and they therefore shift the capability threshold itself. The practical question is which of these levers a government can move before AI rents, contracts, and deployment patterns become institutionally entrenched.

6.1 Make AI Rents Claimable Early

The strongest direct lever is effective public capture, Θ . The practical target is not a single tax instrument or maximal public ownership, but a credible institutional bundle that turns AI rents into durable public claims. Corporate taxation can do part of this work, but only if it is paired with anti-avoidance, withholding, disclosure, and minimum-tax coordination that limit leakage through royalties, transfer pricing, cloud contracts, and intellectual-property location choices (Organisation for Economic Co-operation and Development 2025c,d, 2026). In our model, weak enforcement simply lowers effective Θ and raises the required γ^* .

This implies a concrete administrative agenda. Tax authorities need enough reporting to distinguish domestic AI deployment revenue, cloud and inference payments, model-access royalties, and capital income booked through related entities. Public agencies can support that task by requiring

large AI vendors to disclose contracting structures, compute-region dependencies, and material subcontractors when they sell into public or regulated sectors. The point is to make the AI value chain legible enough that ordinary fiscal tools can operate before rents become hard to locate.

Public-asset channels are the second route. Norway's petroleum fund, Singapore's public-investment model, and Abu Dhabi's sovereign investment institutions are not AI policies, but they show that concentrated rents can be converted into public balance-sheet claims when the state has durable investment and governance capacity (Norwegian Ministry of Finance 2026; Ministry of Finance Singapore 2025; GIC Private Limited 2026; Abu Dhabi Investment Authority 2026; Mubadala Investment Company 2026; ADQ 2026). An AI analogue could take the form of royalties on scarce compute access, public equity in infrastructure, licensing fees, or a social wealth fund capitalized by AI-related revenues.

This is why the result does not imply full-scale nationalization. Moderate, well-governed public claims can move Θ without owning AI firms outright; poorly governed subsidies can raise c or leave Θ effectively unchanged. The key design problem is to make public claims predictable enough for investment, but strong enough that profits from socially consequential automation are not locked into private balance sheets before institutions catch up.

The distinction between statutory and effective capture is crucial for innovation incentives. A levy that deters beneficial investment, induces relocation or avoidance, or is offset by subsidies may raise a headline rate without raising realized Θ . Our model therefore does not identify an optimal tax rate. It shows what a realized public claim would do to fiscal feasibility; policy design must separately test investment, entry, location, and tax-base responses.

6.2 Protect Public Benefit When Making Markets More Competitive

Competition policy and distribution policy need to be designed together. Our oligopoly result is a funding result, not a welfare endorsement: concentration can lower the transfer-funding threshold by preserving rents, while also increasing dependence, exclusion, and lock-in. Competition authorities are therefore right to worry about AI infrastructure and downstream market power (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025b,a). The policy question is what replaces concentrated rents if markets become more contestable.

One answer is stronger public capture. Another is interoperability policy that keeps diffusion broad while still preserving a public claim on the surplus. Public compute investment, open standards, model-access conditions, and public-interest licensing can reduce dependence on a few suppliers without giving up the distributional question. Antitrust success should therefore be measured not only by lower prices or more entrants, but also by whether the resulting market structure leaves institutions capable of financing broad benefit.

This is where the model disciplines a common policy instinct. "More competition" and "more public revenue"

can pull in different directions if competition compresses rents faster than public institutions learn to capture them. A government that weakens gatekeeper power should simultaneously strengthen reporting, interoperability, and public-benefit claims, so diffusion gains are not purchased by making the future rent stream invisible or unclaimable.

6.3 Write Public AI Contracts for Public Benefit

Governments should treat public AI contracting as distribution policy, not only as technology purchasing. Public agencies, schools, courts, health systems, and benefits offices increasingly adopt AI through cloud contracts, model APIs, and application-layer vendors. Those contracts can lower effective deployment costs when they require interoperability, auditability, contestable switching, fallback procedures, and clear responsibility for errors. They can also raise c when they create lock-in, opaque compliance burdens, or off-budget infrastructure costs (Competition and Markets Authority 2024; Organisation for Economic Co-operation and Development 2025a).

Public value has to be visible before a large automation-profit pool exists. Better services, shorter waits, and meaningful error redress can build legitimacy for later public capture; surveillance, brittle automation, and vendor dependence can produce distrust. Procurement is therefore an early opportunity to make AI gains legible as public value rather than private cost cutting.

Public contracts should preserve switching rights, require audit logs for consequential uses, allocate responsibility for errors, limit exclusive data access, and disclose local energy or water burdens. These clauses shape c and the legitimacy of Θ .

6.4 Make Augmentation Benefit Workers

The near-term labor problem is likely to be augmentation before full replacement. Evidence from customer support, knowledge work, and observed model usage suggests that many current systems raise worker productivity or change task mix before they eliminate entire jobs (Brynjolfsson, Li, and Raymond 2023; Dillon et al. 2025; Anthropic 2026). That is why ψ_t matters. In the non-automated block, near-term labor augmentation can be represented as

$$(1 - \bar{\alpha})^{1-\rho} L^\rho \mapsto (1 - \bar{\alpha})^{1-\rho} (\psi_t L)^\rho. \quad (13)$$

This does not replace the automation-rent threshold based on γ_t . It adds a transition channel: AI can raise output and worker productivity before deeper substitution produces the dominant rent pool.

Policy should track wages, task discretion, service quality, and who captures time savings; whether output targets outpace pay; whether training reaches lower-wage workers; and whether workers can contest algorithmic management. Worker voice, sectoral training, support for smaller firms, and public-sector standards are therefore part of AI governance. A positive aggregate ψ_t is not enough if job quality or bargaining power deteriorates. Steering ψ_t toward workers can build legitimacy while institutions prepare to share later automation rents.

This also clarifies the conservative benchmark. Holding the task frontier fixed is not a forecast that new jobs will fail to emerge; it is a stress test for policy if job creation is not enough to absorb displacement.

6.5 Track the Policy Levers Together

The final takeaway is measurement. A capability threshold cannot be governed with a single benchmark score or countdown date. Policymakers need a dashboard linking task-based capability measures to sectoral deployment, profit location, market concentration, public receipts, worker outcomes, and infrastructure costs. METR's task-duration framing is closer to the model's γ_t than decontextualized test scores, but it still needs institutional data about where economic value accrues (Model Evaluation and Threat Research 2025; Arnold et al. 2024).

Responsibility should be explicit. Statistical agencies and independent evaluation bodies should jointly estimate γ_t and $\bar{\alpha}$; tax and finance authorities should publish effective Θ ; competition authorities should monitor θ and ε ; labor agencies should measure ψ_t ; and infrastructure and sector regulators should report c . Each series should include methods, uncertainty intervals, and sectoral dispersion rather than one falsely precise national score.

The dashboard should trigger different responses depending on what moves. Rising capability with weak capture calls for disclosure and enforcement; concentrated profits with declining competition call for interoperability and market governance; visible augmentation with modest profits calls for training, bargaining, and public-service deployment standards. The point is not to find one magic number, but to know whether capability, capture, cost, market structure, or worker benefit is the binding constraint. Regular publication would make the benchmark useful before the threshold is crossed by revealing the binding institutional constraint.

7 Limitations and Policy Directions

Our benchmark is intentionally aggregate. We omit household incidence, endogenous task creation, and strategic profit shifting, though our supplementary Ramsey-Cass-Koopmans analysis (§A.5) shows lower *endogenous* saving relaxes the AI feasibility constraint. These omissions matter for welfare and implementation; our work identifies the *minimal* institutional variables richer models *must* measure.

Key extensions sharpen instrument choice without sacrificing tractability: estimate sector-specific γ_t from observed tasks (Model Evaluation and Threat Research 2025); map Θ to taxes, withholding, royalties, and public assets; place energy, infrastructure, and safety spending inside c ; and model the shift from augmentation ψ_t to substitution. These extensions can turn the dashboard into policy triggers while preserving the result: governance choices materially shift when AI gains become broadly shareable.

Acknowledgements

We thank Leo Kozachkov and Eric Young for feedback and acknowledge support from the Burroughs Wellcome Fund (CASI award) and Foresight Institute.

Ethical Considerations and Distributional Scope

Our fiscal threshold is conditional on a normative premise: some gains from socially consequential automation should be made broadly shareable. The equations do not establish that premise, select a welfare-optimal transfer, or determine who should receive it. UBI is a demanding common yardstick, not a claim that cash transfers always dominate public services, worker ownership, targeted credits, reparative investment, or other uses of public revenue. Democratic institutions still have to justify the objective, instrument, recipients, and tradeoffs.

The U.S. calibration likewise does not imply that the country hosting or taxing a leading AI firm has the sole legitimate claim on its profits. AI value may depend on workers, users, data and content contributors, supply chains, public research, knowledge commons, compute and energy infrastructure, natural resources, and affected communities across many jurisdictions. A nationally collected dollar is therefore a measure of fiscal capacity, not by itself a complete account of distributive justice. Treating globally produced value as belonging only to headquarters countries could deepen existing inequalities in compute access and governance capacity (Lehdonvirta, Wu, and Hawkins 2024; Luna et al. 2024).

Possible responses include source-based withholding, formulaic allocation, royalties, community-benefit agreements, treaty coordination, and internationally governed funds. Our model does not privilege one arrangement. Under any of them, Θ should be interpreted as the share public institutions can direct after legitimate cross-border claims and transfers are recognized. Effective governance also requires representation and accountability for communities supplying labor, data, infrastructure, or resources, rather than assuming that efficient collection alone makes distribution fair.

Finally, neither of the model's two transition channels is ethically self-validating. The lower fiscal threshold under concentration is not an endorsement of monopoly: dependence, exclusion, surveillance, weak contestability, and political power may outweigh the fiscal advantage. Likewise, a positive augmentation effect ψ_t can coexist with work intensification, lower autonomy, or weaker bargaining power. Public reporting, worker voice, contestability, and accessible mechanisms for redress are therefore safeguards for the same transition the benchmark describes, not optional additions after the threshold is crossed.

References

- Abu Dhabi Investment Authority. 2026. Purpose.
- Acemoglu, D.; and Restrepo, P. 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6): 1488–1542.
- ADQ. 2026. Who We Are.
- Aghion, P.; Jones, B. F.; and Jones, C. I. 2017. *Artificial intelligence and economic growth*, volume 23928. National Bureau of Economic Research Cambridge, MA.
- Anderljung, M.; Barnhart, J.; Korinek, A.; Leung, J.; O'Keefe, C.; Whittlestone, J.; Avin, S.; Brundage, M.; Bullock, J.; Cass-Beggs, D.; Chang, B.; Collins, T.; Fist, T.; Hadfield, G.; Hayes, A.; Ho, L.; Hooker, S.; Horvitz, E.; Kolt, N.; Schuett, J.; Shavit, Y.; Siddarth, D.; Trager, R.; and Wolf, K. 2023. Frontier AI Regulation: Managing Emerging Risks to Public Safety. *arXiv:2307.03718*.
- Anthropic. 2026. Anthropic Economic Index Report: Learning Curves.
- Arnold, Z.; Schiff, D. S.; Schiff, K. J.; Love, B.; Melot, J.; Singh, N.; Jenkins, L.; Lin, A.; Pilz, K.; Enweareazu, O.; and Girard, T. 2024. Introducing the AI Governance and Regulatory Archive (AGORA): An Analytic Infrastructure for Navigating the Emerging AI Governance Landscape. *Proceedings of the AAAI/ACM Conference on AI Ethics and Society*, 7(1): 39–48.
- Baumol, W. J. 1967. Macroeconomics of unbalanced growth: the anatomy of urban crisis. *The American economic review*, 57(3): 415–426.
- Brynjolfsson, E.; Li, D.; and Raymond, L. R. 2023. Generative AI at Work. Technical Report 31161, National Bureau of Economic Research.
- Competition and Markets Authority. 2024. AI Foundation Models: Update Paper.
- Congressional Budget Office. 2024. CBO's Economic Forecast: Understanding Productivity Growth. Technical Report 60515, Congressional Budget Office.
- Dafoe, A. 2018. AI Governance: A Research Agenda. Technical report, Centre for the Governance of AI, Future of Humanity Institute, University of Oxford. Version 1.0, August 27, 2018.
- Dillon, E. W.; Jaffe, S.; Immorlica, N.; and Stanton, C. T. 2025. Shifting Work Patterns with Generative AI. Technical Report 33795, National Bureau of Economic Research.
- Erdil, E.; and Besiroglu, T. 2023. Explosive growth from AI automation: A review of the arguments. *arXiv preprint arXiv:2309.11690*.
- GIC Private Limited. 2026. Who We Are.
- Inland Revenue Authority of Singapore. 2026. Corporate Income Tax Rates.
- International Forum of Sovereign Wealth Funds. 2026. Santiago Principles.
- Japan External Trade Organization. 2025. Overview of Corporate Income Taxes (Corporate Tax, Corporate Inhabitant Tax, Enterprise Tax).
- Knoblauch, M.; Roessler, M.; and Zwerschke, P. 2020. The elasticity of substitution between capital and labour in the US economy: A meta-regression analysis. *Oxford Bulletin of Economics and Statistics*, 82(1): 62–82.
- Lehdonvirta, V.; Wu, B.; and Hawkins, Z. 2024. Compute North vs. Compute South: The Uneven Possibilities of Compute-based AI Governance Around the Globe. *Proceedings of the AAAI/ACM Conference on AI Ethics and Society*, 7(1): 828–838.

- Luna, J.; Tan, I.; Xie, X.; and Jiang, L. 2024. Navigating Governance Paradigms: A Cross-Regional Comparative Study of Generative AI Governance Processes & Principles. *Proceedings of the AAAI/ACM Conference on AI Ethics and Society*, 7(1): 917–931.
- Ministry of Finance - United Arab Emirates. 2026. Corporate Tax in the UAE.
- Ministry of Finance Singapore. 2025. Net Investment Returns Contribution.
- Model Evaluation and Threat Research. 2025. Measuring AI Ability to Complete Long Tasks.
- Mubadala Investment Company. 2026. About Mubadala.
- Norwegian Ministry of Finance. 2026. The Government Pension Fund (Statens pensjonsfond).
- O’Keefe, C.; Cihon, P.; Garfinkel, B.; Flynn, C.; Leung, J.; and Dafoe, A. 2020. The Windfall Clause: Distributing the Benefits of AI for the Common Good. In *Proceedings of the 2020 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’20, 327–331. New York, NY, USA: Association for Computing Machinery.
- OpenAI. 2018. OpenAI Charter. Accessed 2026-04-29.
- OpenAI. 2023. Planning for AGI and Beyond. Accessed 2026-04-29.
- Organisation for Economic Co-operation and Development. 2025a. Artificial Intelligence and Competitive Dynamics in Downstream Markets. Technical Report OECD Roundtables on Competition Policy Papers No. 331, OECD.
- Organisation for Economic Co-operation and Development. 2025b. Competition in Artificial Intelligence Infrastructure. Technical Report OECD Roundtables on Competition Policy Papers No. 330, OECD.
- Organisation for Economic Co-operation and Development. 2025c. Corporate Tax Statistics 2025. Technical report, OECD.
- Organisation for Economic Co-operation and Development. 2025d. Withholding Tax Rates and Tax Treaties: Corporate Tax Statistics 2025.
- Organisation for Economic Co-operation and Development. 2026. International Community Agrees Way Forward on Global Minimum Tax Package.
- Sacra. 2025. OpenAI: Revenue, Valuation & Growth Rate.
- Solow, R. M. 1956. A contribution to the theory of economic growth. *The quarterly journal of economics*, 70(1): 65–94.
- Solow, R. M. 1957. Technical change and the aggregate production function. *The review of Economics and Statistics*, 39(3): 312–320.
- Swedish Agency for Economic and Regional Growth. 2025. Taxes and Contributions for Limited Companies.
- The Norwegian Tax Administration. 2025. General Income.
- UBI Center. 2019. Why Not UBI?
- UK Department for Science, Innovation and Technology. 2023. UK Artificial Intelligence Regulation Impact Assessment. Technical report, UK Department for Science, Innovation and Technology.
- U.S. Bureau of Economic Analysis. 2024a. Current-Cost Depreciation of Fixed Assets: Private.
- U.S. Bureau of Economic Analysis. 2024b. Current-Cost Net Stock of Fixed Assets: Private.
- U.S. Bureau of Economic Analysis. 2024c. Gross Domestic Product, Third Quarter 2024 (Advance Estimate).
- U.S. Government Accountability Office. 2022. Corporate Income Tax: Effective Rates Before and After 2017 Law Change. Technical Report GAO-23-105384, U.S. Government Accountability Office.
- World Bank. 2025. Gross Capital Formation (% of GDP) - United States.
- World Economic Forum. 2023. The Future of Jobs Report 2023. Technical report, World Economic Forum.
- Zeira, J. 1998. Workers, machines, and economic growth. *The Quarterly Journal of Economics*, 113(4): 1091–1117.